



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) **EP 1 261 183 A1**

(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
27.11.2002 Bulletin 2002/48

(51) Int Cl.7: **H04L 29/06**

(21) Application number: **01112026.8**

(22) Date of filing: **23.05.2001**

(84) Designated Contracting States:
**AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR**
Designated Extension States:
AL LT LV MK RO SI

• **Ludwig, Reiner**
52393 Hürtgenwald (DE)
• **Sachs, Joachim**
52066 Aachen (DE)

(71) Applicant: **Telefonaktiebolaget L M Ericsson**
(Publ)
126 25 Stockholm (SE)

(74) Representative: **HOFFMANN - EITLE**
Patent- und Rechtsanwälte
Arabellastrasse 4
81925 München (DE)

(72) Inventors:
• **Meyer, Michael**
52062 Aachen (DE)

(54) **Method and system for processing a data unit**

(57) A method of processing a data unit of a first protocol layer (L3) for transmission in a data unit based communication system is described, comprising the steps of passing to a second protocol layer (L2) a given data unit of said first protocol layer (L3) that is to be transmitted, said second protocol layer (L2) lying below said first protocol layer (L3), determining a numeric value of a numerically quantifiable parameter associated with said given data unit of said first protocol layer (L3), embedding said given data unit of said first protocol layer (L3) into one or more data units of said second protocol layer (L2), performing transmission control for said one or more data units of said second protocol layer (L2) that embed said given data unit of said first protocol layer (L3), where said embedding and/or said transmission control is performed in accordance with said numeric value of said numerically quantifiable parameter.

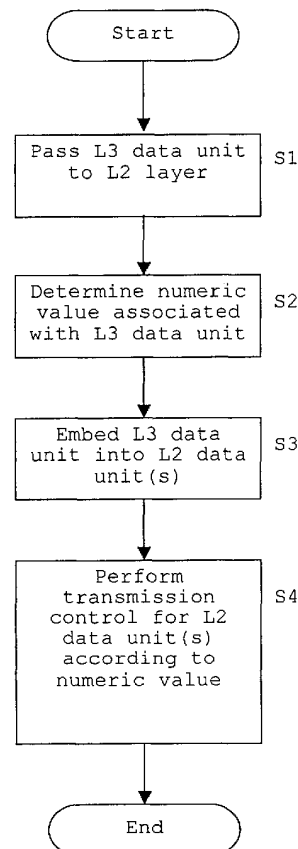


Fig. 1

Description

[Field of the invention]

[0001] The present invention relates to a method of processing a data unit in a data unit based communication system, and to a corresponding data unit based communication system.

[Background of the invention]

[0002] A well known principle for data exchange in networks is that of data unit exchange. This means that data to be sent is broken down into individual units. Rules for sending and receiving such units, as well as rules for the structure of the units themselves, are determined by so-called protocols. Protocols are sets of rules that allow the communication between a sending end and a receiving end, as the rules specify how and in what form data to be sent has to be prepared such that the receiving end may interpret the data and react in accordance to protocol defined rules to which both partners in the communication adhere. The two ends of a communication adhering to a specific protocol are also referred to as peers.

[0003] Such data units are sometimes referred to by different names, depending on the type of protocol involved, such as packets, frames, segments, datagrams, etc. For the purpose of clarity the present description uses the term "data unit" generically for any type of data unit associated with any type of protocol.

[0004] An important concept in communications using data unit exchange is that of protocol layering. This means that a number of protocols (sometimes also referred to as a suite) is organised in a hierarchy of layers, where each layer has specific functions and responsibilities. The concept of layering is well known in the art and described in many textbooks, for example "TCP-IP illustrated, volume 1, The Protocols" by W. Richard Stevens, Addison Wesley, 1994, such that a detailed description is not necessary here.

[0005] The TCP/IP protocol suite is an example of a layered protocol hierarchy. A basic structure of a protocol hierarchy is defined by the OSI (Open System Interconnection) layer model. At a lowest layer, which is also referred to as the physical layer or L1, the functions of directly transporting data over a physical connection are handled. Above the physical layer, a second layer L2, which is also referred to as the link layer is provided. The link layer L2 fulfils the function of handling the transport of data units over links between communication nodes. Above the link layer L2 a third layer L3 is provided, which is also referred to as the network layer. The network layer handles the routing of data units in a given network. An example of a network layer protocol is the internet protocol (IP). Above the network layer, a fourth layer L4 is provided, which is also referred to as the transport layer. Examples of a transport layer protocol

are the transmission control protocol (TCP) or the user datagram protocol (UDP).

[0006] In a data unit based communication system using a hierarchy of protocol layers, a communication comprises passing a given data unit downwards through the protocol hierarchy on the sending side, and passing a data unit upwards through the protocol hierarchy on the receiving side. When a data unit is passed downwards, each protocol will typically perform a certain function with respect thereto, e.g. add further information and change or adapt the structure to specific rules of that protocol layer. Typically each protocol layer will add its own header to a data unit received from a higher protocol layer and may also add delimiters. When a specific protocol layer receives a data unit from a higher protocol layer, it will embed the higher layer data unit into a data unit adhering to the rules of the given protocol layer. The term "embedding" shall refer to both encapsulation in which one data unit of a higher layer is placed into one data unit of a given layer, and to segmentation, where one data unit of a higher layer is segmented into a plurality of data units of the given protocol layer.

[0007] An important aspect of the layering scheme is that the different layers are "transparent". This means that the peers in a layer are oblivious to what happens in another layer.

[0008] Typically, each protocol layer will perform some type of transmission control for its data units. Such transmission control can e.g. comprise the performing of a certain type of forward error correction, the setting of parameters associated with an automatic repeat request (ARQ) function, the scheduling of data units, or the performing of comparable operations.

[0009] It is known to implement protocol layers in such a way that they can be operated in a specific mode with respect to the transmission control. As an example, a so-called numbered mode (or I-mode) and a so-called unnumbered mode (UI-mode) are known. In the numbered mode, if it is determined that a sent data unit was not correctly received by the receiving peer, then the sending peer performs retransmission of said data unit. In this way it can be assured that all packets are correctly transmitted, although this may increase the delay, depending on how many packets have to be transmitted. On the other hand, in the unnumbered mode, no retransmissions are provided. This has the advantage of less delay, but the transmission reliability depends on the quality of the physical connection.

[0010] The possibility of setting a given protocol layer implementation into a specific transmission control mode has the advantage that the selection of the mode can e.g. be performed by a control procedure from a higher protocol layer, in order to optimise the sending of data units from said higher protocol layer. However, it does not provide very much flexibility, as a given protocol layer will typically handle a variety of different types of data units, that require different control settings with respect to the optimisation of the sending of the given

type of data unit. As an example, if an application layer is sending a computer file, it desires to ensure a reliable transmission, and may therefore want to set a lower layer protocol implementation into the numbered mode, or the application layer may want to send data that requires real-time transmission, such as a video stream belonging to a video telephone, in which case transmission speed is more important than reliability, such that the application layer may want to set a lower layer protocol implementation into the unnumbered mode. However, if both computer file data and video data are being sent, then the setting of the lower layer protocol implementation into a given transmission control mode will not provide an optimum solution.

[0011] EP-0 973 302 A1 addresses this problem and proposes a system, in which a given protocol layer that receives higher data unit layers and embeds these higher data layers into data units of said given layer, is arranged to discriminate the higher layer data unit layers by reading the header information and determining the type of the higher layer data unit. Then, a classification is performed in accordance with the identified type. In this way, the given protocol layer can flexibly set the transmission reliability of its data units depending on the type of the higher layer data unit that is embedded in its own data units. As an example, if the system of EP- 0 973 02 A1 is applied to a link layer in the TCP/IP suite, then the link layer can identify if the network layer IP data unit that it receives carries a TCP data unit, in which case the link layer data units are sent in the numbered mode, or if the network layer IP data unit carries UDP data unit, in which case the link layer data units are sent in the unnumbered mode.

[0012] However, the system of EP-0 973 302 A1 is not always practical, as it requires the parsing of higher layer data units in order to identify type information, which e.g. does not work when the higher layer data unit has an encrypted header and/or payload.

[Object of the present invention]

[0013] It is desirable to provide an improved method and system of processing data units of a higher protocol layer at a given protocol layer, which is simple to implement, but flexibly provides improved transmission properties.

[Summary of the invention]

[0014] This object is achieved by a method of processing data units as described in claim 1, and by a data unit based communication system as described in claim 20. Advantageous embodiments are described in the dependent claims.

[0015] In accordance with the present invention, at a given protocol layer, e.g. a link layer L2, one or more numeric values of one or more numerically quantifiable parameters associated with a received given data unit

of a higher protocol layer, e.g. the network layer L3, are determined. In other words, one value of one numerically quantifiable parameter can be determined, or several values of one numerically quantifiable parameter, or one or more respective values for each of a plurality of numerically quantifiable parameters. The at least one numeric value is not derived from information contained in the given data unit of the higher protocol layer. In other words, instead of analysing the content of the higher layer data unit, one or more simple physical properties that are numerically evaluable are measured, and the embedding and/or transmission control is performed in accordance with the determined value. Consequently, no parsing of higher layer data units or other similar complicated processing is necessary.

[0016] As an example, a numerically quantifiable parameter can be the size of the higher layer data unit. In other words, at a given protocol layer e.g. L2, a higher layer data unit, e.g. from the L3 layer, is received, and the size of said L3 data unit is measured. Then the embedding operation for embedding said L3 data unit into one or more L2 data units, or the transmission control operation for transmitting the one or more L2 data units into which said L3 data unit has been embedded, is performed in accordance with the result of said size measurement.

[0017] Preferably, the size measurement is used as a basis for adjusting the transmission control to optimise predetermined target properties. More specifically, the L2 data units are transmitted with parameters set for optimising throughput if the L3 data unit falls into a predetermined size range, and otherwise the L2 data units are transmitted with optimised delay. Namely, if the L3 data unit is found to have a size indicative of a maximum size e.g. falls into a range around or is equal to the TCP maximum segment size if the L3 data units carry TCP data units, then the transmission of the L2 data units, into which said L3 data unit has been embedded, is optimised for throughput, as it may be assumed that the maximum size L3 data unit belongs to a larger amount of data being sent from above the L3 layer. On the other hand, if the L3 data unit is smaller in size, then the transmission control is optimised for delay, as it can be assumed that the smaller L3 data units are associated with control operations, such as synchronisation or acknowledgment messages, where delay optimisation is more suitable than throughput optimisation.

[0018] Other examples of a numerically quantifiable parameter that can be used in the context of the present invention are a buffer fill level of a buffer holding data units of the upper protocol layer (e.g. L3), or of a buffer holding data units of the protocol layer (e.g. L2) receiving the upper layer data units. Another numerically quantifiable parameter is the inter-arrival time of the upper layer data units i.e. the time that passes between the arrival of two consecutive upper layer data units.

[0019] Each of the given examples of numerically quantifiable parameters can be used alone as a basis

for providing one or more values to be used in performing the embedding and/or transmission control, or can be used together with one or more of the named numerically quantifiable parameters for providing such values to be used in performing the embedding and/or transmission control.

[0020] According to a preferred embodiment, the method of the present application also comprises a step of discriminating a group of data units of the higher protocol layer to which the higher layer data unit belongs. This discrimination can occur on the basis of source information and/or destination information in the higher layer data unit and/or a protocol identifier in the higher layer data unit. Especially, if the present preferred embodiment is applied in the context of TCP/IP, the discrimination can consist in determining the flow to which the higher layer data unit belongs. A flow is defined by the source and destination IP address, the source and destination port number and a protocol identifier.

[0021] The result of said discrimination can be used in the transmission control procedure for the data units into which the higher data unit layer is embedded. More specifically, referring to the above described example in which the size of the higher layer data unit was compared with a reference range or a reference value indicative of a maximum data unit size, it is possible to enhance such a comparison by taking into consideration the type of the higher layer data unit. Namely, the type of the higher layer data unit (or the type of a data unit contained in the higher layer data unit) can be used as a means for determining a reference range or reference value with which to compare the size of the higher layer data unit. As an example, if the higher layer data unit is identified as being of the first type (i.e. belonging to a first predetermined group), then the size of said higher data unit is compared to a first size reference value (or to a first set of size reference values, i.e. a discrete range of values, or to a first continuous range of values), whereas if the higher layer data unit is discriminated as being of a second type (i.e. belonging to a second predetermined group), then the size of said higher layer data unit is compared to a second size reference value (or to a second set of size reference values, or to a second continuous range of values).

[0022] On the other hand, the result of said discriminating step can also be used as a basis for determining the numerically quantifiable parameter of which a numeric value is determined. Namely, the numerically quantifiable parameter can be associated with a buffer fill level of a buffer holding data units of the lower protocol layer, where the numeric value is the number of data units of the lower protocol layer (e.g. L2) in the buffer that embed data units of the higher protocol layers belonging to the discriminated group, e.g. belonging to the discriminated flow. Equally, when using the inter-arrival time of the higher layer data unit as the numerically quantifiable parameter, then the numeric value can be chosen as the inter-arrival time value for those higher

layer data units belonging to the discriminated group. In the latter case, i.e. when determining the inter-arrival time of data units belonging to a group, it is possible to add a further condition, e.g. to determine the inter-arrival time of such data units of the group that fall into a given size category. For example, if the group is a flow and the size category is selected to relate to data units that contain an acknowledgement (i.e. have a minimum size), then it is possible to use the inter-arrival time of the acknowledgment data units in the flow as a basis for the transmission control.

[0023] The transmission control performed in accordance with the numeric value of the numerically quantifiable parameter can consist in any suitable measures such as the adjusting of forward error correction, the adjusting of ARQ settings, or the adjusting of scheduling. Furthermore, if the embedding operation comprise a segmentation of the higher layer data unit into several lower layer data units, then this segmentation operation can be performed in accordance with the numeric value, namely by adjusting the size of the lower layer segments in accordance with the determined value. The transmission control can consist in making adjustments for the data units at the layer receiving the higher layer data units, or at a layer below. For example, if the method of the invention is applied to a link layer L2, then adjustments can be made for the L2 data units, but also for L1 data units.

[0024] As already mentioned, the transmission control can be performed to optimise certain target parameters, depending on the determined numeric value. Examples of such target parameters are the throughput (the amount of payload data transported per unit of time) or the (average) transmission delay.

[0025] According to another preferred embodiment, the transmission control for transmitting the lower layer data units, into which a higher layer data unit has been embedded, comprises a discrimination of the lower layer data units, such that each of the one or more lower layer data units is classified into one of a plurality of predetermined transmission categories on the basis of the discrimination result. This discrimination of lower layer data units can advantageously be combined with the discrimination of higher layer data units in such a way that the lower layer data units embedding a higher layer data unit belonging to a predetermined group are themselves divided into sub-groups associated with the group of the higher layer data unit. For example, if the group to which the higher data unit (e.g. L3 data unit) belongs is a flow, then such a flow can be divided into sub-flows at the lower protocol layer (e.g. L2) that embeds the data units of the flow.

[0026] The method and system of the present invention can be put to practice in any suitable or appropriate way by hardware, by software or by any suitable combination of hardware and software.

[Brief description of figures]

[0027] The present invention shall now be described by referring to detailed embodiments, which are only intended to be illustrative, but not limiting, with reference to the appended drawings, in which:

Fig. 1 shows a flow chart for explaining an embodiment of the present invention;

Fig. 2 shows a flow chart for explaining another embodiment of the present invention, in which the numeric value is determined based on data unit size;

Fig. 3 shows a flow chart for explaining another embodiment of the present invention, in which higher layer data units are additionally discriminated;

Fig. 4 shows a flow chart for explaining another embodiment of the present invention, in which the numeric value is determined based on the discrimination result;

Fig. 5 shows an overview of a protocol structure to which the present invention can be applied; and

Fig. 6 schematically illustrates the concept of flow-splitting in the context of the example shown in Fig. 5.

[Detailed description of embodiments]

[0028] In the following, detailed embodiments of the present invention shall be described in order to give the skilled person a full and complete understanding. However, these embodiments are illustrative and not intended to be limiting, as the scope of the invention is defined by the appended claims.

[0029] The following examples shall be described in the context of TCP/IP. However, it may be noted that the present invention is applicable to any protocol hierarchy in which higher layer data units are embedded into lower layer data units, and where a transmission control is performed for the lower layer data units.

[0030] The following embodiments shall be described in the context of applying the invention to a link layer L2, said link layer embedding L3 data units from the network layer. However, this is only an example, and the present invention can be applied at any level of a protocol hierarchy i.e. also below or above the link layer.

[0031] Fig. 1 shows a flow chart describing a first embodiment of the present invention. As can be seen, in a first step S1 a L3 data unit is passed to the link layer L2 below the layer L3. Then, in step S2, a numeric value of a numerically quantifiable parameter associated with

the L3 data unit is determined. It may be noted that the term "associated with the L3 data unit" is to be understood broadly as relating to any numerically quantifiable parameter derivable from the L3 data unit, where said parameter can relate to the L3 data unit as such, or to a data unit from a higher layer than L3 embedded in the L3 data unit. Then, in step S3, the received L3 data unit is embedded into one or more L2 data units. Finally, in step S4, transmission control for the one or more L2 data units and/or L1 data units that embed the received L3 data unit is performed according to the numeric value determined in step S2.

[0032] Although the example of Fig. 1 shows steps S1 to S4 in a specific sequence, this is only an example and other arrangements of steps S1 to S4 are possible.

[0033] The adjustments performed in step S4 will depend on the specific circumstances and the system under consideration. For example, the transmission control may comprise adjusting a forward error correction for the data units of the L2 protocol layer or for the L1 data units below the L2 protocol layer. The type of forward error correction can be chosen as is desirable or appropriate. For example, the transmission control may comprise adjusting a transmission power and/or a data rate (e.g. by selecting a spreading factor in a CDMA system) over a given link and/or a degree of interleaving. The mentioned link can e.g. be a wireless link.

[0034] If the L2 protocol layer comprises a function of providing automatic retransmission of L2 data units under predetermined conditions, then the transmission control may comprise adjusting said retransmission function.

[0035] If the L2 protocol layer comprises a function for the scheduling of the L2 data units, then the transmission control may comprise adjusting said scheduling.

[0036] In the event that the embedding in step S3 is a segmentation operation, then the embodiment of Fig. 1 can be arranged such that the segmentation operation is performed according to said numeric value. In other words, the segmentation operation in dependence on the numeric value can be an alternative to making the transmission control in step S4 dependent on the numeric value, or the segmentation operation in dependence on the numeric value can be a supplement to making the transmission control in step S4 dependent on the numeric value, i.e. can be provided together with such transmission control.

[0037] The adjusting of said segmentation operation may e.g. comprise adjusting the size of the L2 data units into which a given L3 data unit is segmented.

[0038] Fig. 2 shows a flowchart of another embodiment of the present invention. The same reference signs as in Fig. 1 refer to the same steps in Fig. 2. Namely, first an L3 data unit is passed to the L2 layer in step S1. Then, in step S21, which represents an example of the more general step S2 shown in Fig. 1, the size of the received L3 data unit is measured. In other words, it is determined how much data is contained in the received

L3 data unit, in any suitable dimension, such as bits or bytes.

[0039] The L2 data unit is embedded into one or more L2 units in step S3. Then, the transmission control is performed in accordance with steps S41, S42 and S43, which represent a specific example of the general procedure represented as S4 in Fig. 1. More specifically, step S41 compares the measured size of the L3 data unit with at least one reference size, where said reference size preferably represents a maximum data unit size. Then, if the size of the L3 data unit is equal to the reference size, the transmission control is optimised for throughput (step S42), and if the size of the L3 data unit is different from the reference size, then transmission control is optimised with respect to delay (step S43).

[0040] In the above example, step S41 consists in comparing the measured size with a reference value. Naturally, the measured size can also be compared to a range of values, where said range can consist of a discrete number of reference values, or can be continuous, or can be a combination of discrete and continuous values. Also, in the above example, the reference value (or range) is indicative of a maximum data unit size. However, this is only an example, and the reference value (or range) can be chosen in any suitable or desirable way, depending on the given system and circumstances. For example, the reference value can also be indicative of a minimum data unit size that identifies a data unit carrying an acknowledgment.

[0041] The optimising of transmission control for throughput can e.g. consist in reducing the amount of forward error correction and enabling ARQ for ensuring transmission reliability. On the other hand, optimising transmission control for reducing delay can consist in increasing forward error correction, but disabling ARQ. This shall be explained in more detail in the following in the specific context of TCP/IP.

[0042] Non-real-time applications often run on top of TCP, whereas real-time applications usually run on top of UDP. The main design requirement for link layer (L2) protocols for non-real time applications is throughput optimisation. This is for example achieved by tuning the trade-off between forward error correction and ARQ, so that the overall throughput is optimised. It is possible to obtain higher throughputs by weakening forward error correction and enabling ARQ. Throughput is the quantity of transmitted payload per unit of time. The throughput increases with decreasing forward error correction, as forward error correction increases overhead, which leads to a reduced amount of payload being transmitted.

[0043] In an encryption based forward error correction data is divided into blocks, where each block is the result of an encryption operation that combines the payload with redundancy information that allows a decision on the part of the receiver whether a block has been correctly received and decrypted. Then a block error rate may be defined as the rate of events in which a received block is judged as having an error. It may be noted that

reduced forward error correction leads to an increased block error rate, but it can be shown that throughput can nonetheless be increased. For example, when calculating throughput as the product of data rate multiplied by (1 - block error rate), a data rate of 8 kbit/s and block error rate of 0.01 leads to a throughput of 7.92 kbit/s, whereas a data rate of 12 kbit/s (achieved by less forward error correction) and a block error rate of 0.1 (due to the reduced forward error correction) lead to a throughput of 10.8 kbit/s. Consequently, although the block error rate is ten times higher, the increased payload data rate leads to an overall increased throughput.

[0044] The problem is that even in a flow corresponding to a TCP based bulk data transfer, which should overall be optimised for throughput, there are individual data units or segments, which are delay sensitive. Examples are the connection set-up messages, which should be exchanged as quickly as possible to be able to start with data transmission. Another example is the TCP acknowledgement messages in reverse direction, or HTTP requests, which should be received as soon as possible.

[0045] In order to differentiate between L3 data units that should be optimised for throughput and such that should be optimised for delay reduction, the present embodiment determines the size of the L3 data unit, see step S21 in Fig. 2.

[0046] Namely, L3 data units for which the throughput should be optimised in general belong to a larger application data amount, which is segmented into transport layer data units, e.g. TCP segments. Therefore, these transport layer data units, which are then embedded in network layer (L3) data units generally have a maximum transfer unit size, e.g. the TCP maximum segments size (MSS).

[0047] Consequently, as already mentioned, step S41 of Fig. 2 can be implemented in such a way that the comparison is conducted with respect to a reference value or reference range indicative of a predetermined maximum transfer unit size. Typically, the value of the maximum transfer unit size is 256 byte, 512 byte, 536 byte or 1460 byte, such that taking into account the IP header of 40 byte, suitable reference size values used in step S41 could be 296 byte, 552 byte, 576 byte or 1500 byte.

[0048] As already mentioned, step S41 can be implemented in such a way that the measured size of the L3 data unit is compared with a plurality of reference sizes, e.g. the previously mentioned series of discrete values 296, 552, 576 and 1500 bytes, or with a suitable continuous range, or with a combination of discrete values and continuous values. Then, if the measured size of the L3 unit is equal to any one of the plurality of reference sizes, or falls into the predetermined range or ranges, then step S42 is enabled, to thereby optimise the transmission control for throughput.

[0049] In this connection, it can be mentioned that the reference size values may not only take into account the IP header, but also the possibility of header compres-

sion. In other words, the plurality of size reference values used in step S41 also contains a set of values that take into account header compression.

[0050] In the example of Fig. 2, other data units having different sizes than the reference size are sent with optimised delay (step S43). In other words, they are sent with minimum delay. In the example of TCP/IP this means that all pure TCP control messages (such as synchronisation messages, acknowledgement messages,...) are sent with minimum delay. Also, the final data unit belonging to a larger application data amount is also transmitted with optimised (e. g. minimum) delay.

[0051] As already indicated above, delay optimisation can for example be achieved by choosing the appropriate coding scheme and/or adjusting the transmission power. If one assumes an additional delay for a retransmission to be 100 ms, one can conclude that the transmission delay becomes less and less important for small data units. Therefore, it is better to use stronger forward error correction (e.g. more redundancy) and/or higher transmission power to improve the mean delay.

[0052] For example, if one considers a data unit of 100 byte size, and using the previously mentioned example values of a data rate of 8 kbit/s, block error rate of 0.01 and a block size of 20 byte, then a mean delay of 100 + 5 ms results, or assuming a data rate of 12 kbit/s, a block error rate of 0.1 and a block size of 30 byte, then a mean delay of 80 + 40 ms results. In the first example, one needs five blocks to transmit 100 byte, in the second only 4, due to the increased data rate and therefore block size. In the first case, the probability that the one of the five blocks is corrupted is approximately 5 %, leading on average to an additional delay of 5 ms. The same calculation provides an additional delay of 40 ms for the second case. Consequently, the increased forward error correction, although the data rate is reduced, leads to an improved mean delay.

[0053] Another example is when the transmission control comprises regulation via the transmission power. This can be achieved e.g. by setting a target signal-to-interference ratio for the desired link, which is maintained by a power control loop. A higher signal-to-interference ratio results in an improved link quality, i.e. the number of transmission errors on the link is reduced, at the price of increased transmission power. Less transmission errors reduce the delay of error recovery and therefore the mean transmission delay is reduced.

[0054] In the example of Fig. 2, the outcome of decision step S41 is to either optimise the transmission for throughput or for delay. Naturally, this is only an example. The transmission control can also be optimised with regard to other target parameters, as is suitable or desired for a specific application or system. Such target parameters include system resources like e.g. transmission power, spreading codes, number of assigned physical channels. Furthermore, the outcome of decision step S41 can also lead to more than the two procedures S42, S43. For example, it is possible that step S41 is

implemented in such a way that the L3 data unit size is compared to n reference sizes or n reference ranges, and the outcome leads to a corresponding number of n transmission control procedures, each transmission control procedure corresponding to one size reference value or range, and additionally one further default transmission control procedure for all those L3 data units that do not have the size of one of the n size reference values or does not fall into one of the n ranges.

[0055] In the embodiments described above, the transmission control performed in step S4 or steps S41 to S43 was performed on the basis of the numeric value determined in step S2. It should be noted that the transmission control can also be based on further measurements or determinations conducted with respect to a received L3 data unit. Namely, as shown in figure 3, in which steps S1, S21 and S3 are identical to the corresponding steps in Fig. 2, such that a further description is not necessary, a step S5 is performed after step S3, said step S5 discriminating a group to which the L3 data unit belongs. Then, in step S40, which is an example of general step S4 of Fig. 1, the transmission control for the one or more L2 data units that embed the received L3 data unit, is performed on the basis of the numeric value determined in step S21 and the discrimination result of the discrimination step S5.

[0056] Preferably, the discrimination is performed on the basis of destination information contained in the L3 data unit and/or a protocol identifier contained in the L3 data unit. For example, the group into which the L3 data unit is classified can be the flow to which it belongs. As already mentioned previously, in the context of TCP/IP, the group into which the L3 data unit is classified can be the flow to which it belongs, and the flow is defined by a source and destination IP address, a source and destination port number, and a protocol identifier. As another example, the group to which an L3 data unit belongs can also be determined by the type of payload being transmitted, which type can e.g. be determined by checking a protocol identifier in the data unit. This means checking the protocol identifier of the L3 data unit itself, or checking one or more protocol identifiers of data units being transported as payload in the L3 data unit.

[0057] The performing of the transmission control in step S40 can be conducted in any suitable or desirable way, depending on the specific circumstances of the application. For example, procedure S40 of Fig. 3 can be arranged in a similar way as steps S41, S42 and S43 of Fig. 2, with the addition that the size reference used in the comparison is also selected in dependence on the group into which the L3 data unit is classified. For example, the group can be determined on the basis of the type of data unit embedded in the L3 data unit, e.g. the discrimination consists in determining if the IP data unit carries a TCP data unit or a UDP data unit. If it is determined that it carries a TCP data unit, then a first set of size reference values can be used, and if it is determined that the IP data unit carries a UDP data unit, then a sec-

ond set of size reference values can be used in the comparison.

[0058] In Fig. 3, the discrimination result obtained in step S5 is used as an element in the transmission control operation of procedure S40. However, the discrimination of a received L3 data unit into one of a predetermined number of groups can also be used as a basis for determining the numeric value that is then later to be used in the transmission control operation. This is shown in Fig. 4. After having received the L3 data unit in step S1, the discrimination step S5 is performed. Then, in step S22, which is an example of the general step S2 shown in Fig. 1, the numeric value is determined on the basis of the discrimination result. Thereafter steps S3 and S4 are conducted, as already explained in connection with Fig. 1.

[0059] For example, the numerically quantifiable parameter can be associated with a buffer fill level of a buffer holding L2 data units, where the numerically quantifiable parameter is the number of L2 data units in the buffer that embed L3 data units belonging to the group discriminated in step S5. The numerically quantifiable parameter can alternatively be the inter-arrival time of L3 data units belonging to the discriminated group. The advantages of determining the numerically quantifiable parameter in this way shall be discussed in more detail with respect to the example shown in Fig. 5.

[0060] Fig. 5 shows an embodiment of the present invention implemented in the context of a link layer operating in accordance with the universal mobile telephone system (UMTS) as e.g. described in the Technical Specification 3G TS 25.301 V3.3.0 (1999-12), published by the 3rd Generation Partnership Project (<http://www.3gpp.org>). This document is herewith fully incorporated by reference into the present disclosure.

[0061] It may be noted that Fig. 5 shows a preferred application of the present invention. However, the present invention is by no means restricted thereto. Much rather, the method and system of the present invention can be applied in the context of any mobile communication device, e.g. also one operating in accordance with the general package audio service (GPRS) or any other mobile communication standard. Furthermore, a method and system of the present invention are naturally not restricted to the application in mobile communication systems, but can be applied in any communication system having a protocol hierarchy.

[0062] In Fig. 5, a radio resource control (RRC) procedure 201 is provided, which controls the operation of other parts of the general L2 implementation, said L2 protocol implementation consisting of sub-layers L2/PDCP, L2/BMC, L2/RLC and L2/MAC. Fig. 5 schematically shows control connections 204, 205, 206, 207 and 208 to the other procedures, that shall be described in the following. Namely, a one or more packet data convergence protocol (PDCP) procedures 202 are implemented for performing the functions of the packet data convergence protocol. Furthermore, a BMC (Broadcast/

Multicast Control) procedure 203 is implemented for performing the function of controlling the sending of data units to a plurality of destinations, via multicasting or broadcasting. Reference numeral 209 refers to radio link control (RLC) procedures that implement the radio link control protocol. 211 identifies the medium access control (MAC) part of the L2 layer. Finally, 302 schematically represents the physical layer L1.

[0063] The radio resource control procedure 201 receives control information from the higher layer L3 (not shown). The user-plane passes down L3 data units that are to be embedded into L2 data units, to the PDCP procedure 202, to the BMC procedure 203 and/or a RLC procedure 209.

[0064] Reference numeral 210 represents logical channels in the L2 layer, while reference numerals 301 refers to transport channels between L2 and L1.

[0065] The general control for the L2 layer implementation is performed by the RRC procedure 201. More specifically, user data is transmitted on a radio access bearer. Via the radio resource control procedure 201, one or more radio bearers are configured, which determine the layer 1 and layer 2 configuration of the radio protocols. In the PDCP procedure 202, IP header compression is applied and a multiplexing of multiple traffic flows onto one logical channel is used. The RLC procedure applies backward error correction, such as an ARQ, for a logical channel, among other things. The MAC procedure 211 performs a scheduling of L2 data units and performs the distribution onto the transport channels 301.

[0066] As an embodiment of the present invention, the function of the PDCP procedure 202 is extended by a splitting function, which separates the received traffic flow (e. g. an IP flow) into sub-flows, which are then transmitted via separate RLC connections. Since IP header compression is applied in PDCP, the necessary flow information is already available in PDCP. An example, a corresponding flow splitting is schematically depicted in Fig. 6. A flow of IP data units received at the PDCP procedure 202 is split into four sub-flows, each respectively handled by a separate RLC procedure 209a, 209b, 209c and 209d.

[0067] The splitting is performed in accordance with the numeric value of the numerically quantifiable parameter, e.g. the L3 data unit size, a buffer fill level, or an inter-arrival time. Additionally, the splitting can also be conducted in dependence on the type of data unit.

[0068] For example, the numerical value can be the L3 unit size and the type of data unit being transported. Then the splitting shown in Fig. 6 can be arranged in such a way that data units having the size of an acknowledgment and being TCP data units can be transmitted with an optimum RLC configuration, i.e. minimum delay, and the other IP data units are separated onto RLC connections, which are optimised for the respective data unit size and/or respective data unit type.

[0069] An alternative way of splitting the IP flow into a plurality of RLC sub-flows is to separate different phas-

es of the IP flow, e.g. to separate the TCP slow-start phase from the congestion avoidance phase. If the slow-start phase is configured such that a lower transmission delay is achieved, the end-to-end performance will be greatly enhanced and link resources (e.g. codes) are released more quickly.

[0070] Methods for reducing the transmission delays are e.g. a higher signal/interference-ratio target, stronger forward error correction, more aggressive ARQ settings etc. To differentiate the different phases of a TCP flow, the above described methods of employing a buffer fill-level as the numerically quantifiable parameter, or using the inter-arrival time can be used. As already mentioned previously, the determination of an inter-arrival time can also be combined with the determination of a data unit size range, in that the value used as a basis for transmission control is the inter-arrival time of data units falling into a predetermined size range.

[0071] For example, if the RLC buffer level is used, then it could be concluded that if the RLC buffer fill level is below a predetermined threshold, the TCP flow is probably in the slow-start phase (or at the end of application data), since no new data is arriving. Above this predetermined threshold, or above a second predetermined threshold, it is very likely that TCP is in the congestion avoidance phase. As a consequence, if the buffer fill level indicates a slow-start phase, then the L2 data units are placed into an RLC connection optimised for reducing delay, and if the buffer fill level indicates a congestion avoidance phase, then the L2 data units are placed into a RLC connection optimised for throughput.

[0072] Corresponding considerations apply when using the inter-arrival time of the IP data units as a numerically quantifiable parameter. Namely, if the inter-arrival time is above a predetermined threshold, then this may indicate a slow-start phase, whereas if the inter-arrival time is below this same predetermined threshold, or below a second threshold, then this may indicate a congestion avoidance phase.

[0073] It may be noted that an alternative to splitting the flow into several RLC sub-flows, as shown in Fig. 6, in which each stream of data units is handled by a separate RLC procedure 209a, 209b, 209c or 209d, is to transmit all data units in one adaptive sub-flow. In this case, all data units are handled by the same RLC procedure, however, each data unit is handled differently (different forward error correction, different segmentation sizes, different ARQ, different transmission power (e.g. different target signal-to-interference ratio), etc.).

[0074] Although the present invention has been described with reference to specific embodiments these embodiments only serve to provide the skilled person with a full and complete understanding of the invention, but are not intended to limit the scope. The scope of the invention is much rather defined by the appended claims. Reference numerals in the claims are intended to make the claims easier to understand and also do not restrict the scope.

Claims

1. A method of processing a data unit of a first protocol layer (L3) for transmission in a data unit based communication system, comprising the steps of:

passing to a second protocol layer (L2) a given data unit of said first protocol layer (L3) that is to be transmitted, said second protocol layer (L2) lying below said first protocol layer (L3);

determining one or more numeric values, said one or more numeric values belonging to at least one numerically quantifiable parameter associated with said given data unit of said first protocol layer (L3);

embedding said given data unit of said first protocol layer (L3) into one or more data units of said second protocol layer (L2),

performing transmission control for said one or more data units of said second protocol layer (L2) that embed said given data unit of said first protocol layer (L3),

where said embedding and/or said transmission control is performed in accordance with said one or more numeric values of said at least one numerically quantifiable parameter.

2. The method of claim 1, wherein said at least one numerically quantifiable parameter is the size of said given data unit of said first protocol layer (L3).
3. The method of claim 1 or 2, wherein said at least one numerically quantifiable parameter is associated with a buffer fill level of a buffer holding data units of said first protocol layer (L3) or said second protocol layer (L2).
4. The method of one of claims 1 to 3, wherein said at least one numerically quantifiable parameter is an inter-arrival time of data units of said first protocol layer (L3).
5. The method of one of claims 1 to 4, furthermore comprising a step of discriminating a group of data units of said first protocol layer (L3) to which said given data unit of said first protocol layer (L3) belongs.
6. The method of claim 5, wherein said group is discriminated on the basis of source information and/or destination information and/or a protocol identifier contained in said data units of said first protocol layer (L3).

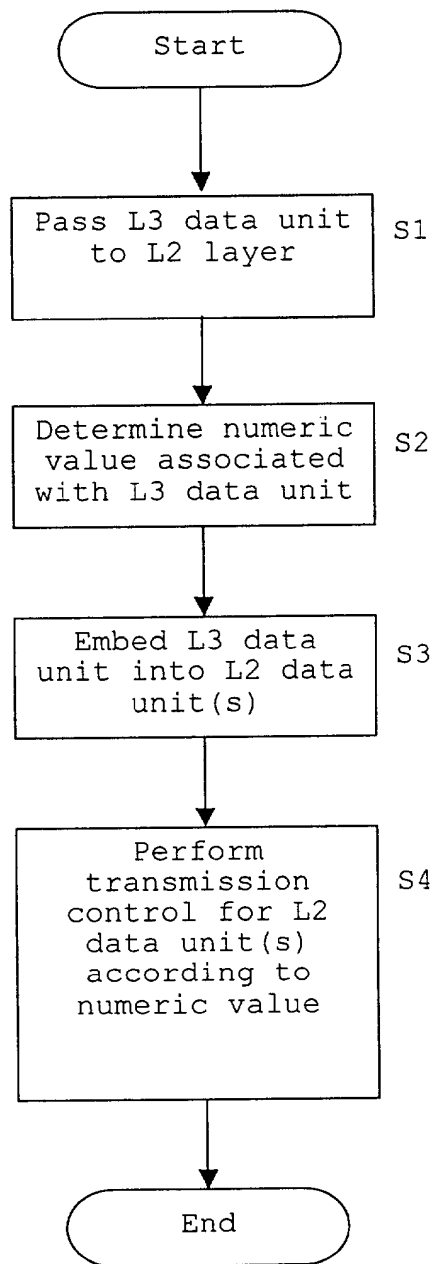
7. The method of claim 5 or 6, wherein said transmission control for said one or more data units of said second protocol layer (L2) that embed said given data unit of said first protocol layer (L3) is also performed in accordance with a result of said discriminating step. 5
8. The method of claim 5 or 6, wherein said numeric value is determined on the basis of a result of said discriminating step. 10
9. The method of claim 8, wherein said numerically quantifiable parameter is associated with a buffer fill level of a buffer holding data units of said second protocol layer (L2), said numerically quantifiable parameter being the number of data units of said second protocol layer (L2) in said buffer that embed data units of said first protocol layer (L3) belonging to said group. 15
10. The method of claim 8, wherein said numerically quantifiable parameter is the inter-arrival time of data units of said first protocol layer (L3) belonging to said group. 20
11. The method of one of claims 1 to 10, wherein said transmission control comprises adjusting a forward error correction for said data units of said second protocol layer (L2) or for data units of a third protocol layer (L1) below said second protocol layer (L2). 25
12. The method of claim 11, wherein a function (RLC) controls the sending of said data units of said second protocol layer (L2) over a link, and said transmission control comprises adjusting a transmission power and/or a data rate over said wireless link and/or a degree of interleaving. 30
13. The method of one of claims 1 to 12, wherein said second protocol layer (L2) comprises a function of providing automatic retransmission of data units of said second protocol layer (L2) under predetermined conditions, and where said transmission control comprises adjusting said retransmission function. 35
14. The method of one of claims 1 to 13, wherein said second protocol layer (L2) comprises a function (MAC) for scheduling of said data units of said second protocol layer (L2), and where said transmission control comprises adjusting said scheduling. 40
15. The method of one of claims 1 to 14, wherein said embedding comprises a segmentation operation for data units of said first protocol layer (L3), and where said transmission control comprises adjusting said segmentation operation. 45
16. The method of claim 15, wherein the adjusting of said segmentation operation comprises adjusting the size of the data units of said second protocol layer (L2) into which said given data unit of said first protocol layer (L3) is segmented. 50
17. The method of one of claims 1 to 16, wherein said transmission control comprises discriminating said one or more data units of said second protocol layer (L2) in which said given data unit of said first protocol layer (L3) is embedded on the basis of said numeric value, and placing each of said one or more data units of said second protocol layer (L2) into one of a plurality of predetermined transmission categories on the basis of said discrimination result. 55
18. A computer program arranged to execute the method of one of claims 1 to 17.
19. A computer program storage medium storing a computer program according to claim 18.
20. A data unit based communication system comprising implementations of a first protocol layer (L3) and a second protocol layer (L2), said first protocol layer (L3) lying above said second protocol layer (L2), said system being arranged to pass a given data unit of said first protocol layer (L3) that is to be transmitted to said second protocol layer (L2), and said implementation of said second protocol layer (L2) being arranged to determine one or more numeric values belonging to at least one numerically quantifiable parameter associated with said given data unit of said first protocol layer (L3), embed said given data unit of said first protocol layer (L3) into one or more data units of said second protocol layer (L2), and perform transmission control for said one or more data units of said second protocol layer (L2) that embed said given data unit of said first protocol layer (L3) in accordance with said one or more numeric values of said at least one numerically quantifiable parameter.
21. The data unit based communication system of claim 20, wherein said at least one numerically quantifiable parameter is the size of said given data unit of said first protocol layer (L3).
22. The data unit based communication system of claim 20 or 21, wherein said at least one numerically quantifiable parameter is associated with a buffer fill level of a buffer holding data units of said first protocol layer (L3) or said second protocol layer (L2).
23. The data unit based communication system of one of claims 20 to 22, wherein said at least one numerically quantifiable parameter is an inter-arrival time

of data units of said first protocol layer (L3).

24. The data unit based communication system of one of claims 20 to 23, wherein said implementation of said second protocol layer (L2) is arranged to discriminate a group of data units of said first protocol layer (L3) to which said given data unit of said first protocol layer (L3) belongs.
25. The data unit based communication system of claim 24, wherein said implementation of said second protocol layer (L2) is arranged to discriminate said group on the basis of source information and/or destination information and/or a protocol identifier contained in said data units of said first protocol layer (L3).
26. The data unit based communication system of claim 24 or 25, wherein said implementation of said second protocol layer (L2) is arranged to also perform said transmission control for said one or more data units of said second protocol layer (L2) that embed said given data unit of said first protocol layer (L3) in accordance with a result of said discriminating.
27. The data unit based communication system of claim 24 or 25, wherein said implementation of said second protocol layer (L2) is arranged to determine said numeric value on the basis of a result of said discriminating step.
28. The data unit based communication system of claim 27, comprising a buffer for holding data units of said second protocol layer (L2), said numerically quantifiable parameter being the number of data units of said second protocol layer (L2) in said buffer that embed data units of said first protocol layer (L3) belonging to said group.
29. The data unit based communication system of claim 27, comprising a timer for measuring an inter-arrival time of data units of said first protocol layer (L3) belonging to said group, where said numerically quantifiable parameter is the inter-arrival time of data units of said first protocol layer (L3) belonging to said group.
30. The data unit based communication system of one of claims 20 to 29, wherein said transmission control comprises adjusting a forward error correction for said data units of said second protocol layer (L2) or for data units of a third protocol layer (L1) below said second protocol layer (L2).
31. The data unit based communication system of claim 30, further comprising a function (RLC) for controlling the sending of said data units of said second protocol layer (L2) over a link, and said transmission

control comprises adjusting a transmission power and/or a data rate over said link and/or a degree of interleaving.

32. The data unit based communication system of one of claims 20 to 31, wherein said implementation of said second protocol layer (L2) comprises a function of providing automatic retransmission of data units of said second protocol layer (L2) under predetermined conditions, and where said transmission control comprises adjusting said retransmission function.
33. The data unit based communication system of one of claims 20 to 32, wherein said implementation of said second protocol layer (L2) comprises a function (MAC) for scheduling of said data units of said second protocol layer (L2), and where said transmission control comprises adjusting said scheduling.
34. The data unit based communication system of one of claims 20 to 33, wherein said implementation of said second protocol layer (L2) is arranged to perform a segmentation operation for data units of said first protocol layer (L3), and where said transmission control comprises adjusting said segmentation operation.
35. The data unit based communication system of claim 34, wherein the adjusting of said segmentation operation comprises adjusting the size of the data units of said second protocol layer (L2) into which said given data unit of said first protocol layer (L3) is segmented.
36. The data unit based communication system of one of claims 20 to 35, wherein said transmission control comprises discriminating said one or more data units of said second protocol layer (L2) in which said given data unit of said first protocol layer (L3) is embedded on the basis of said numeric value, and placing each of said one or more data units of said second protocol layer (L2) into one of a plurality of predetermined transmission categories on the basis of said discrimination result.
37. The data unit based communication system of one of claims 20 to 35, wherein said system is contained in a mobile communication device.
38. The data unit based communication system of claim 37, wherein said mobile communication device operates in accordance with the standard of the Universal Mobile Telecommunication System.

**Fig. 1**

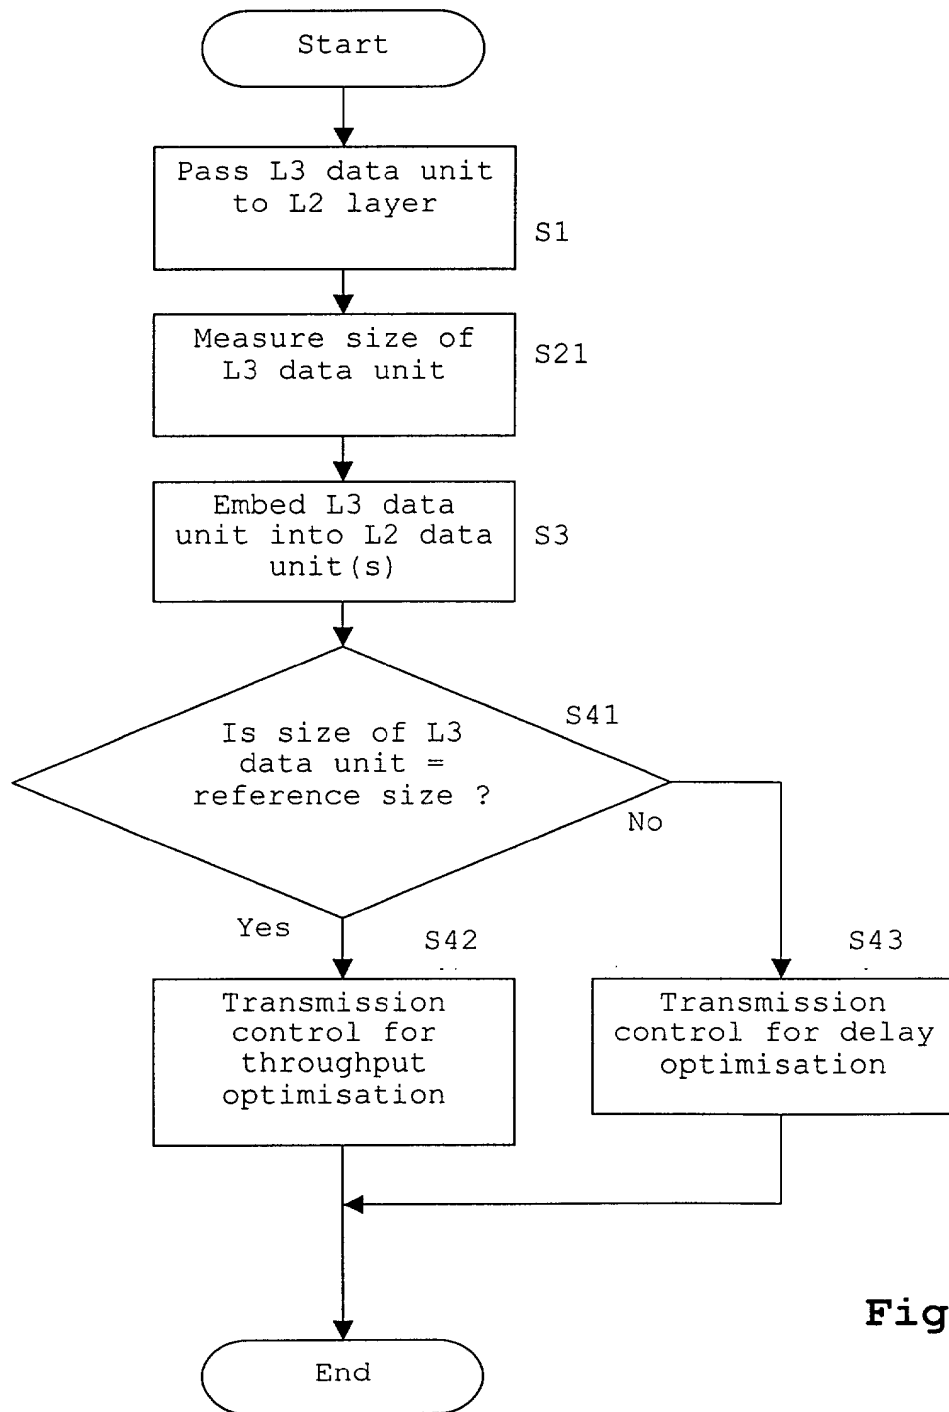


Fig. 2

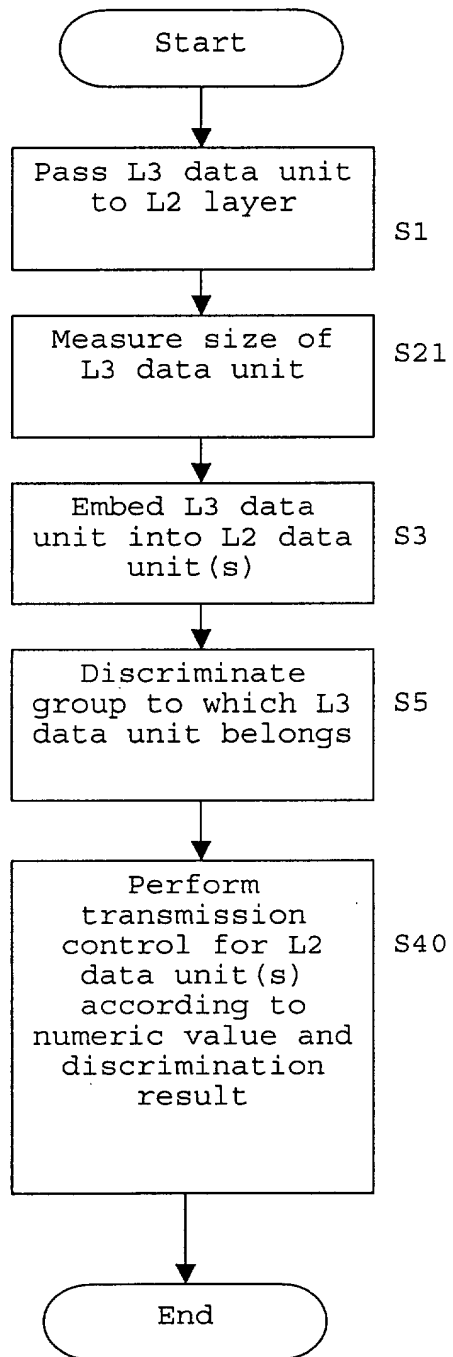


Fig. 3

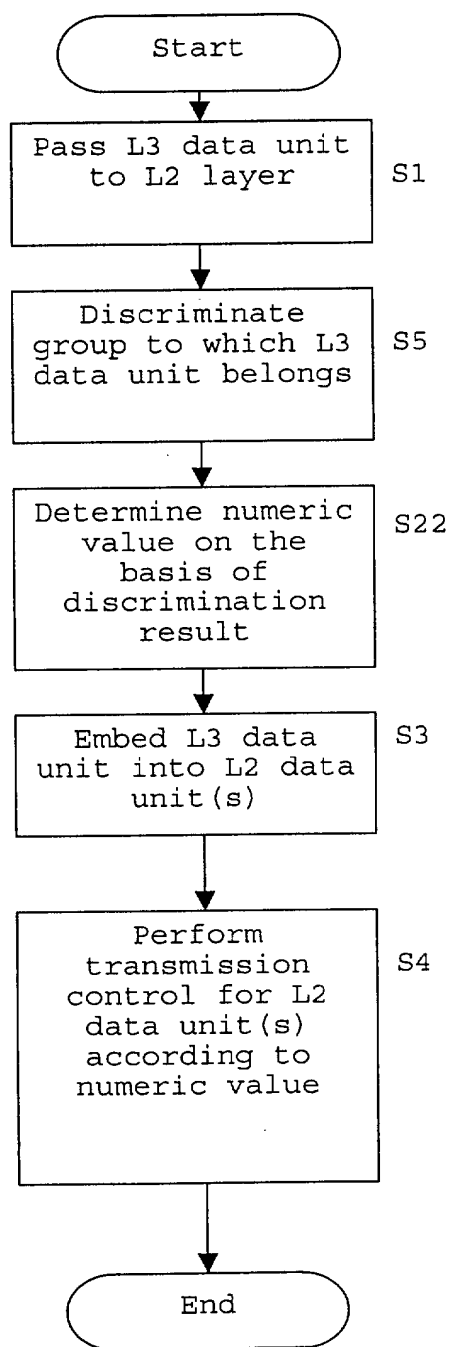



Fig. 4

Fig. 5






Fig. 6



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 01 11 2026

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	EP 0 959 589 A (LUCENT TECHNOLOGIES INC) 24 November 1999 (1999-11-24) * column 3, line 27 - column 6, line 37 * * column 8, line 31 - column 9, line 58 * * column 13, line 7 - line 22 * * column 23, line 50 - column 24, line 37 *	1,3, 13-15, 18-20, 22, 32-34, 37,38	H04L29/06
A		2,4,5, 11,17, 21,23, 24,30,36	
X	US 5 701 300 A (DO MI SUN ET AL) 23 December 1997 (1997-12-23) * column 1, line 27 - column 2, line 16 *	1,2, 18-21	
A	EP 0 905 950 A (SONY CORP ;NIPPON TELEGRAPH & TELEPHONE (JP)) 31 March 1999 (1999-03-31) * column 3, line 30 - line 57 * * column 5, line 9 - line 48 *	3,22	TECHNICAL FIELDS SEARCHED (Int.Cl.7) H04L
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 25 September 2001	Examiner Brichau, G
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document	

EPO FORM 1503 03.82 (PB4001)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 01 11 2026

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

25-09-2001

Patent document cited in search report		Publication date	Patent family member(s)		Publication date
EP 0959589	A	24-11-1999	AU	2906599 A	02-12-1999
			BR	9912983 A	24-07-2001
			CN	1280448 A	17-01-2001
			CN	1268854 A	04-10-2000
			EP	0964548 A1	15-12-1999
			EP	0959589 A2	24-11-1999
			JP	2000049866 A	18-02-2000
			JP	2000083077 A	21-03-2000
US 5701300	A	23-12-1997	KR	132959 B1	21-04-1998
EP 0905950	A	31-03-1999	JP	11098219 A	09-04-1999
			CN	1218346 A	02-06-1999
			EP	0905950 A1	31-03-1999
			SG	74080 A1	18-07-2000
			US	6233224 B1	15-05-2001

EPO FORM P4458

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82